arXiv:1602.01376v1 [cs.NA] 3 Feb 2016

INV-ASKIT: A Parallel Fast Direct Solver for Kernel Matrices

Chenhan D. Yu*, William B. March†, Bo Xiao‡ and George Biros§
*Department of Computer Science
*†‡§ Institute for Computational Engineering and Science
The University of Texas at Austin, Austin, Texas, USA
*chenhan@cs.utexas.edu, †march@ices.utexas.edu, ‡bo@ices.utexas.edu, †gbiros@acm.org

We present a parallel algorithm for computing the approximate factorization of an N-by-N kernel matrix. Once this factorization has been constructed (with $N \log^2 N$ work), we can solve linear systems with this matrix with $N \log N$ work. Kernel matrices represent pairwise interactions of points in metric spaces. They appear in machine learning, approximation theory, and computational physics. Kernel matrices are typically dense (matrix multiplication scales quadratically with N) and ill-conditioned (solves can require 100s of Krylov iterations). Thus, fast algorithms for matrix multiplication and factorization are critical for scalability.

Recently we introduced ASKIT, a new method for approximating a kernel matrix that resembles N-body methods. Here we introduce INV-ASKIT, a factorization scheme based on ASKIT. We describe the new method, derive complexity estimates, and conduct an empirical study of its accuracy and scalability. We report results on real-world datasets including "COVTYPE" (0.5M points in 54 dimensions), "SUSY" (4.5M points in 8 dimensions) and "MNIST" (2M points in 784 dimensions) using shared and distributed memory parallelism. In our largest run we approximately factorize a dense matrix of size 32M $\times$ 32M (generated from points in 64 dimensions) on 4,096 Sandy-Bridge cores. To our knowledge these results improve the state of the art by several orders of magnitude.

I. INTRODUCTION

Given a set of N points $x_i \in \mathbb{R}^d$ and a kernel function $\mathcal{K}(x_i, x_j) : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$, a kernel matrix is the $N \times N$ matrix[1] whose entries are given by $K_{ij} = \mathcal{K}(x_i, x_j)$ for $i, j = 1, \ldots, N$. We are interested in the factorization of K.

Kernel matrices are typically dense, so matrix-vector multiplication with K (henceforth, "MATVEC") requires $\mathcal{O}(N^2)$ work. Solving a linear system with K (which we will refer to as "INVERT" K) is not stable because K can be nearly singular. We typically invert $\lambda I + K$, where λ is the *regularization parameter* and I is the $N \times N$ identity matrix. The LU factorization of $\lambda I + K$ requires $\mathcal{O}(N^3)$ work. An iterative Krylov method requires $\mathcal{O}(N^2)$ work but the constant can be quite high depending on λ.

Motivation and significance. Kernel matrices appear in N-body problems [11], machine learning methods for classification, regression, and density estimation [9], [16], [30], spatial statistics [7], Gaussian processes [27], and approximation theory [28], among other applications. Typical operations required with kernel matrices are matrix-vector multiplication, solving linear systems, and eigenvalue computations. For example, in machine learning, solving linear systems is used in ridge regression and support vector machines. Eigenvalue computations are required in spectral clustering. Such operations can be done either with iterative solvers in which we only require matrix-vector multiplication. However, the conditioning can be bad and convergence can be quite slow.

For many (but not all) kernel matrices, dramatic acceleration of the MATVEC and the INVERT operations can be achieved if we exploit subtle properties of their structure. For example, consider the Gaussian kernel, a popular kernel in machine learning,

$$\mathcal{K}(x_i, x_j) = \exp\left(-\frac{1}{2}\frac{\|x_i - x_j\|_2^2}{h^2}\right), \qquad (1)$$

where h is the *bandwidth*. For small h, K approaches the identity matrix whereas for large h, K approaches the rank-one constant matrix. The first regime suggests sparse approximations while the second regime suggests global (numerical) low-rank approximations. For the majority of h values, however, K is neither sparse nor globally low-rank. It turns out that the *off-diagonal blocks* of K may admit good low-rank approximations. A key challenge is identifying and constructing these approximations. When the points are in low dimensions (say $d < 10$), effective techniques for the MATVEC and INVERT operations exist. But for high-dimensional problems, there is very little work.

Contributions: Our goal is a method that exhibits algorithmic and parallel scalability for inverting $\lambda I + K$ for large N and d. We only require the ability to compute an entry of K in $\mathcal{O}(d)$ time. Our scheme extends ASKIT[2], a method we introduced in a series of papers [22], [23], [25]. ASKIT

[1] We restrict our discussion to the case in which x_i and x_j come from the same point set. The i-indexed points are called the target points, and the j-indexed points are called the source points. A kernel matrix can also involve target and source points that come from different sets.

[2] *A*pproximate *S*keletonization *K*ernel *I*ndependent *T*reecode. The ASKIT library is available at http://padas.ices.utexas.edu/libaskit.

approximates K in $\mathcal{O}(dN \log N)$ time. Here we introduce INV-ASKIT for factorizing and inverting the ASKIT-based approximation of K. In particular:

- We present sequential and parallel factorization algorithms that have $\mathcal{O}(N \log^2 N)$ complexity and can be applied in $\mathcal{O}(N \log N)$ time (described in §II-C, §II-D).
- We present a theoretical analysis of the complexity of the algorithm (in terms of time, communication and storage). The estimates for the parallel factorization and solution are given by (16) and (15) respectively in §II-D.
- An ASKIT variant approximates K as a triple sum: a block-diagonal plus a low-rank plus a *sparse* matrix. In this case we use INV-ASKIT as a preconditioner for a GMRES iterative solver.
- We apply our solver to a binary classification problem using kernel regression. We present the solver's accuracy, training and test errors, and compare it with Krylov iterative methods (described in §IV).
- We conduct strong and weak scaling analysis in which we report communication, computation and overall efficiency (described in §IV).

Our method uses the hierarchical decomposition of ASKIT. The fundamental concept behind it is simple: First, approximate K as the sum of a block-diagonal matrix and a low-rank matrix. Second, compute the inverse using the Sherman-Morrison-Woodbury (SMW) formula. Third, apply this decomposition recursively to each block of the block-diagonal matrix. *To our knowledge,* INV-ASKIT *is the only algorithm that can be used with high-dimensional data for a large variety of kernels and the only algorithm that supports shared and distributed memory parallelism.* INV-ASKIT *does not assume symmetry of the kernel, global low-rank, sparsity, or any other property other than that, up to a small sparse matrix-correction, the off-diagonal blocks admit a low-rank approximation.*

Related work: We do not review the large literature of accelerating the matrix-vector operation with a kernel matrix. We focus only on fast factorization schemes. Many times kernel matrices admit a good *global low-rank* approximation. For such matrices the *Nystrom method* and its variants can be very effective [8], [14], [31] for both approximating K and its regularized inverse. However, in many applications K does not have a global low-rank structure and Nystrom methods fail. In [24], we discussed this in detail and presented classification results on real datasets for which the Nystrom method fails.

The idea of exploiting the block-diagonal-plus-low-rank structure has been discussed extensively in the past fifteen years for kernel matrices in low dimensions and finite-difference/finite element matrices. Indeed for point datasets in low dimensions, there exist methods that are much faster and more scalable than ASKIT. Approximating $(\lambda I + K)^{-1}$ is more challenging, but this is also a well-studied problem. Efficient schemes for approximating the inverse of kernel (but also other types of) matrices are also known as *fast direct solvers*, and they have roots in nested dissection (or domain decomposition), multifrontal methods and truncated Newton-Schulz iterations. Examples include [1]–[3], [5], [10], [13], [15], [20], [29] and this list is non-exhaustive. All these works, however, have been applied successfully only in low-dimensional problems ($d < 4$) and do not address parallelism and performance.

To our knowledge, [18] is the earliest work that parallelizes hierarchical matrices with bulk synchronous parallelization on shared memory systems, but it is not efficient due to the $\mathcal{O}(N)$ critical path. In [19], $\mathcal{H}$-LU improves its efficiency by employing a DAG-based task scheduling. Distributed implementations include [29] and [17], but, to our understanding, they were designed for sparse matrices related to 2D and 3D stencil discretizations, not kernel matrices. Although the basic machinery is similar, but the matrices are quite different. We discuss this more in §II-A. For Gaussian kernels, the only work for fast inversion is [2], [3] in which the authors consider Gaussian and other radial-basis kernels in low dimensions ($d < 4$). Parallelism and scalability were not considered in that work. Their largest test had 1M points in 3D. We consider datasets that have $30\times$ more points and over $200\times$ higher dimensions.

Limitations: This is a novel algorithm and this paper represents our first attempt in an efficient implementation. For this reason it has several limitations. First, we do not provide a theoretical analysis of its numerical stability. (But it is relatively easy to prove it for symmetric and positive definite matrices.) Second, we do not exploit certain symmetries in ASKIT that can further reduce the factorization time. For example, K can be symmetric. Third, we do not performance-tune various algorithm parameters. Fourth, we do not discuss optimizations that can be used to amortize the factorization cost over different values of λ. Fifth, for certain ASKIT variants our scheme can be only used as a preconditioner and the errors can be significant. Sixth, there are additional low-rank structures that we do not exploit. Finally, there exist fundamental limitations inherited from ASKIT. Not all kernel matrices have a block-diagonal-plus-low-rank structure. Examples are point datasets that have high *ambient dimensionality* and kernels that are highly oscillatory (e.g., the high-frequency Helmholtz kernel). Also, the sparse-matrix correction for the off-diagonals can have large norm, in which case using this version of INV-ASKIT as a preconditioner will not work.

II. METHODS

Notation: In Table I, we summarize the main notation used in the paper. We begin with discussing the basic concepts in factorization of hierarchical matrices §II-A. Then we present a brief summary of the ASKIT algorithm §II-B.

Notation	Description
N	problem size
d	dimension of input points
m	leaf node size
s	number of skeleton points
p	number of MPI processes
α, β	tree nodes, overloaded with the indices they own
$\mathtt{l}$, $\mathtt{r}$	left and right children of the treenode
$\mathcal{X}$	$\mathbb{R}^{d\times N}$, point dataset
$\mathcal{X}(\alpha)$	points owned by node α, i.e. $\mathcal{X}(\alpha) = \{x_i \vert \forall i \in \alpha\}$
$\widetilde{K}$	approximate kernel matrix

Table I *Main notation used. Also occasionally we will use MATLAB-style matrix representations.*

We conclude with the sequential §II-C and parallel §II-D factorization algorithms.

A. Hierarchical Matrices and Treecodes

Informally, $K \in \mathbb{R}^{N\times N}$ is hierarchical if it can be approximated by $D + UV$, where D is a ***block-diagonal matrix***, UV is ***low-rank***, and ***each block of*** D ***is also a hierarchical matrix***. To define the D, U, V matrices first we consider the following block partitioning to $N/2$-by-$N/2$ blocks

$$K = \begin{bmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{bmatrix} = \begin{bmatrix} K_{11} & 0 \\ 0 & K_{22} \end{bmatrix} + \begin{bmatrix} 0 & K_{12} \\ K_{21} & 0 \end{bmatrix}. \quad (2)$$

D is the first matrix in the sum and the product of UV is equal to the second matrix; U and V will be defined shortly. So in summary: (1) Blocks K_{12} and K_{21} admit low-rank *approximations*, which implies that U and V are low-rank. (2) The same decomposition applies recursively for the diagonal blocks K_{11} and K_{22}. This approximation derived by recursive partitioning is directly associated with a binary partitioning of the rows and columns of K. Note that using lexicographic ordering of the rows like in (2) will not, in general, result in low-rank off-diagonal blocks. An appropriate reordering, in our case using the tree structure, is critical for exposing such low-rank structure. Given the finest granularity m, we partition the points using geometric information and a binary tree so that we have less than m points per leaf. (The number of points m is manually selected. Typically is about 1048 or 2048 and it effects both computation time and accuracy. For a discussion, see [22]). We denote the levels of the tree by l, with $l = 0$ at the root and $l = \log_2(N/m)$ at the leaves. Let α be a node in the tree. We "overload" the notation by using α to also indicate the indices of the points assigned to α so that all $\mathcal{X}(\alpha) \in \alpha$. Then the hierarchical partition and approximation of K assumes that for every node α, we can partition the self-interactions between its points as follows

$$K_{\alpha\alpha} = \begin{bmatrix} K_{\mathtt{ll}} & K_{\mathtt{lr}} \\ K_{\mathtt{rl}} & K_{\mathtt{rr}} \end{bmatrix} = \begin{bmatrix} K_{\mathtt{ll}} & \\ & K_{\mathtt{rr}} \end{bmatrix} + \begin{bmatrix} & K_{\mathtt{lr}} \\ K_{\mathtt{rl}} & \end{bmatrix}, \quad (3)$$

where $\mathtt{l}$ denotes the left child of a node, and $\mathtt{r}$ the right child of a node. Equivalently to (2), we can write $\widetilde{K}_{\alpha\alpha} = D_\alpha + U_\alpha V_\alpha$. The matrices U_α and V_α are selected to approximate off-diagonal blocks in (3) such that

$$\begin{bmatrix} 0 & K_{\mathtt{lr}} \\ K_{\mathtt{rl}} & 0 \end{bmatrix} \approx U_\alpha V_\alpha. \quad (4)$$

The particular form of these matrices for the ASKIT method is given in (4). If K is invertible, we can write immediately the ***approximate inverse*** of K by using the SMW formula:

$$\begin{aligned} \widetilde{K} &= D + UV = D(I + D^{-1}UV) \\ &= D(I + WV), \text{ where } W = D^{-1}U. \\ \widetilde{K}^{-1} &= (I + WV)^{-1}D^{-1} \\ &= (I - W(I + VW)^{-1}V)D^{-1} \\ &= (I - WZV)D^{-1}, \text{ where } Z = (I + VW)^{-1}. \end{aligned} \quad (5)$$

The extension to $\lambda I + K$ is trivial. In INV-ASKIT, factorizing K amounts to visiting every tree node α to compute W and Z. We require traversing the descendants of α to compute $D^{-1}U$. If s is the numerical rank of $K_{\mathtt{lr}}$ and $K_{\mathtt{rl}}$, then computing W requires $2s$ linear solves with D_α. Computing Z requires factorizing (using LAPACK) the $2s \times 2s$ matrix $I + VW$. The recursion stops at the leaves of the tree in which we have a single matrix $D_\alpha = K_{\alpha\alpha}$, which we then again factorize using LAPACK. We give pseudocode for the precise algorithms in §II-C.

As mentioned in §I, there is extensive work on fast solvers for hierarchical matrices. Such solvers rely on the SMW formula. But the details can differ significantly, for example the permutation or reordering of the matrix, construction of U and V, and the exploitation of various symmetries.[3] ASKIT can also produce a slightly different decomposition (known as a fast multipole matrix or an $\mathcal{H}$-matrix), which decomposes K into a sum of a hierarchical matrix $D + UV$ (which has the form we discussed above) and a ***sparse matrix*** S. Next, we give some details for ASKIT's construction of the hierarchical approximation of K and briefly explain the $\widetilde{K} = D + UV + S$ decomposition.

B. ASKIT

Details on the algorithm can be found in [22]. Here we briefly summarize it. The input is a set of points $\mathcal{X}$ and a list of nearest neighbor points in $\mathcal{X}$ for each point $x_i \in \mathcal{X}$. We build a balanced binary tree-based partitioning of the points. We use a parallel ball tree [26] data structure. For $d < 4$ quadtrees and octrees can be effectively used, but in high dimensions such data structures cannot be used.

After building the tree, we construct the hierarchical representation. In ASKIT we use the notion of a point being *separated* from a tree node α. A basic separation criterion is that $i \notin \alpha$. If $K_{i\alpha}$ is the interaction between i and α, we

[3] More precisely our matrix can be described as a HODLR (Hierarchical Off-Diagonal Low Rank) matrix [1]. Other types of matrices include hierarchically semi-separable matrices (HSS), and $\mathcal{H}$-matrices [12].

construct $K_{i\tilde{\alpha}}$ and $P_{\tilde{\alpha}\alpha}$ such that

$$K_{i\alpha} \approx K_{i\tilde{\alpha}} P_{\tilde{\alpha}\alpha}. \tag{6}$$

The $\tilde{\alpha} \subset \alpha$ points are referred to as the "***skeletons***". The skeletons are the pivot columns of the QR factorization of a subsample of the block matrix $K_{[1:N]\alpha}$. (The nearest-neighbor information is used to effectively sample this block matrix.) $K_{i\tilde{\alpha}} = \mathcal{K}(x_i, \mathcal{X}(\alpha))$ does not involve any approximation, but instead of computing interactions with all points in α we just compute interactions with $\tilde{\alpha}$. $P_{\tilde{\alpha}\alpha}$ is computed using the interpolative decomposition algorithm. A key property of ASKIT is that the skeletons of a node are selected from the skeletons of its children. That is, for a node α the skeletons $\tilde{\alpha}$ are a subset of $\tilde{\mathtt{l}} \cup \tilde{\mathtt{r}}$ or $[\tilde{\mathtt{l}}\tilde{\mathtt{r}}]$. Furthermore $P_{\tilde{\alpha}\alpha} w(\alpha)$ (when we compute Kw) can be computed recursively by only storing $P_{\tilde{\alpha}[\tilde{\mathtt{l}}\tilde{\mathtt{r}}]}$. If we assume the skeleton size s is fixed for all nodes, then $P_{\tilde{\alpha}[\tilde{\mathtt{l}}\tilde{\mathtt{r}}]}$ is an $s \times 2s$ matrix.

The details of constructing skeletons and the P matrices can be found in [22]. In the pseudocode below we outline the two steps required to perform a fast `MatVec` $u = Kw$. The "*Skeletonize*" step builds the projection matrices P for each node and computes $w_{\tilde{\alpha}} = P_{\tilde{\alpha}\alpha} w_\alpha$, auxiliary variables that are then used in the "*Evaluate*" step in which for every target point i, the tree is traversed top-down to compute u_i the entry of the output vector for point i.

```
1:  //Skeletonize:
2:  Skeletonize(α)
3:  if α is leaf then
4:      Find α̃
5:      Construct P_{α̃α}
6:      w_α̃ = P_{α̃α} w_α
7:  else
8:      Skeletonize(l)
9:      Skeletonize(r)
10:     Construct α̃, P_{α̃[l̃r̃]}
11:     w_α̃ = P_{α̃[l̃r̃]} w_{[l̃r̃]}
12: end if
```

```
1:  //Evaluate:
2:  MatVec (i, α)
3:  if Separated(i, α) then
4:      u_i = K_{iα̃} w_α̃
5:  else
6:      if α is leaf then
7:          u(i)+ = K_{iα} w_α
8:      else
9:          MatVec(i, l)
10:         MatVec(i, r)
11:     end if
12: end if
```

In the *Skeletonize* algorithm, lines 4, 5 and 10 do not depend on w. Therefore, if we want to apply K to a different vector w, only lines 6 and 11 are needed. In the *Evaluate* phase, ASKIT provides two different algorithms for the `Separated`(i, α) function. In the first algorithm, i is separated from α if $i \notin \alpha$. This definition creates a hierarchical approximation of K identical to (3). Then, if $\mathtt{l}$ and $\mathtt{r}$ are two sibling nodes with parent α, the matrices $K_{\mathtt{lr}}$ and $K_{\mathtt{rl}}$ are approximated by

$$K_{\mathtt{lr}} \approx K_{\mathtt{l}\tilde{\mathtt{r}}} P_{\tilde{\mathtt{r}}\mathtt{r}} \text{ and } K_{\mathtt{rl}} \approx K_{\mathtt{r}\tilde{\mathtt{l}}} P_{\tilde{\mathtt{l}}\mathtt{l}}. \tag{7}$$

We now can define for a node α

$$U_\alpha = \begin{bmatrix} K_{\mathtt{l}\tilde{\mathtt{r}}} & 0 \\ 0 & K_{\mathtt{r}\tilde{\mathtt{l}}} \end{bmatrix} \text{ and } V_\alpha = \begin{bmatrix} 0 & P_{\tilde{\mathtt{r}}\mathtt{r}} \\ P_{\tilde{\mathtt{l}}\mathtt{l}} & 0 \end{bmatrix}. \tag{8}$$

This completes the definition of the hierarchical matrix for INV-ASKIT (compare (8) with (4)). To reduce memory costs however, ASKIT never constructs $P_{\tilde{\alpha}\alpha}$ explicitly. Instead, looking at the *Skeletonize* function, ASKIT only stores $P_{\tilde{\alpha}[\tilde{\mathtt{l}}\tilde{\mathtt{r}}]}$ for each node α. In the factorization we will need the whole $P_{\tilde{\alpha}\alpha}$ matrix. We explain how to compute it in §II-C.

Other then $i \notin \alpha$, ASKIT can also define i to be separated from α if no *nearest-neighbor* of i is in α. This definition introduces a more complicated structure in the matrix. We do construct (7) but for points i that have *nearest-neighbors* in α, line 4 in the *Evaluate* is replaced by $u_i = K_{i\alpha} w_\alpha$, i.e., an exact evaluation. For certain kernels this significantly increases the accuracy of the overall scheme. A convenient way to think about this is that the second algorithm builds an approximation that is equivalent to

$$K \approx (D + UV) + S = \widetilde{K} + S, \tag{9}$$

where $D + UV$ is the hierarchical matrix and S is a sparse matrix that is zero if we use the $i \notin \alpha$ separation criterion. If the norm of S is small relative to $\widetilde{K}$ we expect that $\widetilde{K}$ will be a good preconditioner for $\widetilde{K} + S$. (Again there are kernels for which this is not the case.)

In this section we summarized the basic ideas in ASKIT. Its structure is very similar to what is known as a *treecode* in computational physics. We omitted several important details of the ASKIT algorithm. For example, we assumed that the number of skeleton points s is fixed. In reality, this number differs for each tree node and it is chosen adaptively so that the approximation error is below a tolerance τ. Finally, the parallelization and scalability of ASKIT was described in [25]. Another version of ASKIT that resembles the relationship between fast-multipole methods and treecodes is described in [23]. In these papers, we also discuss the error analysis and work/storage complexity of ASKIT. Next, we will assume that we are given $D + UV$ from ASKIT and we proceed with discussing the INV-ASKIT factorization in detail using the ideas we outlined in §II-A.

C. Fast direct solver

Given the binary tree, the regularization parameter λ, and $\widetilde{K}$ in terms of $\tilde{\alpha}$ and $P_{\tilde{\alpha}\alpha}$ for every tree node α, we construct a factorization of $\lambda I + \widetilde{K}$ that allows the fast solution of linear systems. We describe the algorithm for the case of $\lambda = 0$, for notational simplicity, but the algorithms are stated for any λ. So given $u \in \mathbb{R}^d$, we wish to compute $w = \widetilde{K}^{-1} u$.

In the factorization phase we compute U, W, V and Z. We first show how (5) is computed in a blocked 2×2 matrix. $D(I + WV)$ is computed as

$$\begin{bmatrix} \widetilde{K}_{\mathtt{ll}} & \\ & \widetilde{K}_{\mathtt{rr}} \end{bmatrix} \left(I + \begin{bmatrix} W_{\mathtt{l}\tilde{\mathtt{r}}} & \\ & W_{\mathtt{r}\tilde{\mathtt{l}}} \end{bmatrix} \begin{bmatrix} & P_{\tilde{\mathtt{r}}\mathtt{r}} \\ P_{\tilde{\mathtt{l}}\mathtt{l}} & \end{bmatrix} \right). \tag{10}$$

where $W_{\mathtt{l}\tilde{\mathtt{r}}} = \widetilde{K}_{\mathtt{ll}}^{-1} K_{\mathtt{l}\tilde{\mathtt{r}}}$ and $W_{\mathtt{r}\tilde{\mathtt{l}}} = \widetilde{K}_{\mathtt{rr}}^{-1} K_{\mathtt{r}\tilde{\mathtt{l}}}$. Then the

partitioned form of $I - WZV$ is

$$I - \begin{bmatrix} W_{\mathtt{l}\tilde{\mathtt{r}}} & \\ & W_{\mathtt{r}\tilde{\mathtt{l}}} \end{bmatrix} \begin{bmatrix} I & P_{\tilde{\mathtt{r}}\mathtt{r}} W_{\mathtt{r}\tilde{\mathtt{l}}} \\ P_{\tilde{\mathtt{l}}\mathtt{l}} W_{\mathtt{l}\tilde{\mathtt{r}}} & I \end{bmatrix}^{-1} \begin{bmatrix} & P_{\tilde{\mathtt{r}}\mathtt{r}} \\ P_{\tilde{\mathtt{l}}\mathtt{l}} & \end{bmatrix}. \tag{11}$$

If α is an non-leaf node, $P_{\tilde{\alpha}\alpha}$ is evaluated recursively by

$$P_{\tilde{\alpha}\alpha} = P_{\alpha[\tilde{\mathtt{l}}\tilde{\mathtt{r}}]} \begin{bmatrix} P_{\tilde{\mathtt{l}}\mathtt{l}} & \\ & P_{\tilde{\mathtt{r}}\mathtt{r}} \end{bmatrix} \text{ using } \texttt{GEMM}. \tag{12}$$

At the leaf nodes $P_{\tilde{\alpha}\alpha}$ is computed by ASKIT (using the LAPACK routines `GEQP3` and `GELS` [22]).

Algorithm II.1 `Factorize`(α)

if α is leaf **then**
  LU factorization $\lambda I + K_{\alpha\alpha}$. /* $\mathcal{O}(m^3)$ */
else
  `Factorize`(`l`) and `Factorize`(`r`).
  Compute $K_{\mathtt{l}\tilde{\mathtt{r}}}$, $K_{\mathtt{r}\tilde{\mathtt{l}}}$, $P_{\tilde{\mathtt{l}}\mathtt{l}}$, $P_{\tilde{\mathtt{r}}\mathtt{r}}$. /* $\mathcal{O}(s^3 + Ns^2)$ */
  $W_{\mathtt{l}\tilde{\mathtt{r}}} = \texttt{Solve}(\mathtt{l}, K_{\mathtt{l}\tilde{\mathtt{r}}})$ and $W_{\mathtt{r}\tilde{\mathtt{l}}} = \texttt{Solve}(\mathtt{r}, K_{\mathtt{r}\tilde{\mathtt{l}}})$.
  LU factorize the reduced system in (11). /* $\mathcal{O}(s^3)$ */
end if

Factorization Phase. Algorithm II.1 performs a postorder traversal of the tree. If α is a leaf node, we factorize $\lambda I + K_{\alpha\alpha}$ with partial-pivoting LU factorization (LAPACK routine `GETRF`). Otherwise, we first compute $K_{\mathtt{l}\tilde{\mathtt{r}}}$, $K_{\mathtt{r}\tilde{\mathtt{l}}}$, $P_{\tilde{\mathtt{l}}\mathtt{l}}$, $P_{\tilde{\mathtt{r}}\mathtt{r}}$ using (12). Then we solve $W_{\mathtt{l}\tilde{\mathtt{r}}} = \widetilde{K}_{\mathtt{l}\mathtt{l}}^{-1} K_{\mathtt{l}\tilde{\mathtt{r}}}$ and $W_{\mathtt{r}\tilde{\mathtt{l}}} = \widetilde{K}_{\mathtt{r}\mathtt{r}}^{-1} K_{\mathtt{r}\tilde{\mathtt{l}}}$ by calling `Solve()` (Algorithm II.2). (Note that when we visit a node at some level, the factorization information required for the solve is already in place.) We compute and factorize Z using the blocked representation in (11). `Solve()` is described in Algorithm II.2, which applies $\widetilde{K}_{\alpha\alpha}^{-1}$ to a vector u.

Algorithm II.2 $w = \texttt{Solve}(\alpha, u)$

if α is leaf **then**
  LU solver $w = (\lambda I + K_{\alpha\alpha})^{-1} u$. /* $\mathcal{O}(m^2)$ */
else
  $v_{\mathtt{l}} = \texttt{Solve}(\mathtt{l}, u_{\mathtt{l}})$ and $v_{\mathtt{r}} = \texttt{Solve}(\mathtt{r}, u_{\mathtt{r}})$.
  $v = [v_{\mathtt{l}}\ v_{\mathtt{r}}]$
  Compute $w = v - WZVv$ using (11). /* $\mathcal{O}(sN)$ */
end if

Solving Phase. If α is a leaf node, $w = K_{\alpha\alpha}^{-1} u$ directly invokes an LU solver (LAPACK routine `GETRS`). Otherwise, we first solve $v_{\mathtt{l}} = \widetilde{K}_{\mathtt{l}\mathtt{l}}^{-1} u_{\mathtt{l}}$ and $v_{\mathtt{r}} = \widetilde{K}_{\mathtt{r}\mathtt{r}}^{-1} u_{\mathtt{r}}$ recursively. Then we compute (11) with two matrix-matrix multiplications (BLAS routine `GEMM`) and an LU solver (`GETRS`). Note that both algorithms are implemented using bottom-up tree traversal and "for- loops"—we do not actually employ recursion.

Refactoring for different λ. In the case that we want to refactor for a different value of λ, (e.g., in cross-validation studies), $\widetilde{K}$ needs to be refactorized, but the skeletons and V we have computed remain the same. Thus, when we refactor $\widetilde{K}$ with different λ, only W and Z need to be recomputed. (The factorization at the leaf level can be also amortized if SVD is used instead of LU, but we do not discuss this here.)

Complexity of the sequential algorithm. We present the complexity analysis of Algorithm II.1 and Algorithm II.2. The main parameters are s and m. They depend on the kernel, the dimensionality, and the geometric distribution of points. They control the accuracy of the scheme, but are not explicitly related to problem size N. In the actual implementation s is chosen adaptively for each node. If we use its maximum value (across nodes), then the results below are an upper bound to the work complexity. Note that in some applications, the kernel does depend on the problem size, for example the bandwidth value for the Gaussian kernel in machine learning problems. In that case m and s will need to change accordingly. But here we assume a fixed kernel for all values of N, since the way the kernel changes with increasing N depends on the application and the geometry of the dataset.

$T^f(N)$ denotes the complexity of the factorization Algorithm II.1 with N points, and $T^s(N)$ denotes the complexity of Algorithm II.2 with N points. We can verify that

$$T^s(N) = 2T^s\left(\frac{N}{2}\right) + \mathcal{O}(Ns) = \mathcal{O}(sN \log N). \tag{13}$$

Using the solving phase complexity, we derive $T^f(N)$ as

$$T^f(N) = 2T^f\left(\frac{N}{2}\right) + s^2 T^s(N) = \mathcal{O}(s^2 N \log^2 N). \tag{14}$$

To summarize, Algorithm II.1 takes $\mathcal{O}(N \log^2 N)$ work and Algorithm II.2 takes $\mathcal{O}(N \log N)$ work.

Storage. Beyond the storage required for ASKIT, we need to store U (overwritten by W), V and the LU factorization in each level. U and V take $\mathcal{O}(2sN)$ in each level. Overall, the space requirement is $\mathcal{O}(dN + 2sN \log \frac{N}{m})$. To save about half of the memory, V can be constructed recursively when it's needed, and we discard it when we reach the parent level. This will result in slightly larger constants in the solving phase.

Shared memory parallelism: There are several ways to exploit shared memory parallelism. We have implemented and tested two of those. The first way is level-by-level parallel tree traversal for both construction and evaluation, in which we use sequential versions of BLAS/LAPACK routines, concurrent traversal of nodes at the same level, and sequential traversal across levels. The second way is to use sequential traversals and use the shared-memory parallelism for the BLAS/LAPACK routines. It turns out, that it is more efficient to use a hybrid version of these two approaches. From level-4 to the leaf level, we disable parallel BLAS and LAPACK because we have experimentally observed that parallel tree traversal provides higher efficiency due to better memory locality. From level-0 to level-3 in the local

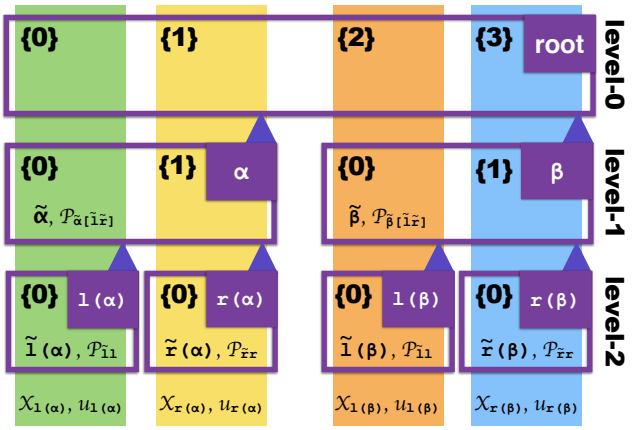

Figure 1 *The* ASKIT *distributed binary tree with $p = 4$. The purple boxes indicate tree nodes. Depending on the level a tree node is logically distributed across several MPI processes, which are organized under an MPI-communicator. The ranks are indicated in curly brackets. Each process owns the data of its colored rectangle region. For example, the yellow process owns $\mathtt{r}(\alpha)$ and $P_{\tilde{\mathtt{r}}\tilde{\mathtt{r}}}$ at level 2, but it doesn't own $\tilde{\alpha}$ at level 1. We depict the main data used in the factorization, the matrices P and the skeleton points at each node. Notice that the nodes at level 2 are not the leaf nodes. But the computations in the subtrees at higher levels are entirely local to the MPI processes that own them. In this example we have four such processes, indicated by different colors.*

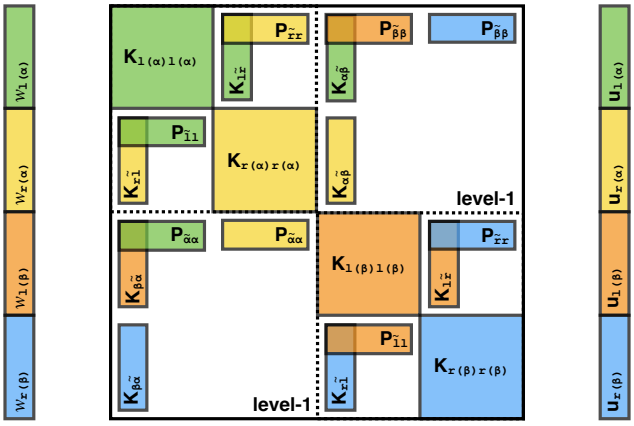

Figure 2 *Distributed $K_{\alpha\tilde{\beta}}$ and $P_{\alpha\tilde{\alpha}}$ factors of the $\widetilde{K}$ matrix that corresponds to the partitioning of Figure 1. Each MPI processes (in the global communicator) stores only the matrices that have the same color with it. A matrix that has multiple color is distributed across several processes. We also show the partitioning of input and output vectors.*

tree, we switch from parallel tree traversal to parallel BLAS operations. In parallel tree traversal some cores may start to idle from tree level three, which only has 8 tree nodes. So in the top four levels of the tree we use parallel `GETRF`, `GETRS` and `GEMM`.

D. Parallel algorithm

In [25] we discussed the parallelization of ASKIT. Here we build on that work and derive our parallel factorization algorithm. We assume that we have p ranks in the global MPI communicator. The data-structure used in ASKIT is summarized in Figure 1. The distributed-memory parallelism in the algorithm is invoked between tree levels-0 and -$\log p$. We briefly review the structure of the tree before discussing the algorithms and the complexity.

Distributed Tree. Without loss of generality, we assume N and p are powers of two and that each MPI process owns $n = N/p$ points. Figure 1 is a distributed tree with $p = 4$. We use purple boxes to indicate tree nodes. Each treenode node is assigned to a number of MPI processes. All processes sharing a node are grouped to an MPI communicator. We use $\{\mathrm{i}\}$ to indicate the i_{th} MPI rank in the local communicator. For example, at level 0 the root of the tree is distributed among all 4 processes (from green to blue in Figure 1). The green process has local MPI rank $\{0\}$, and the blue process has rank $\{3\}$. If a node α contains q processes, then the processes with rank $\{i < \frac{q}{2}\}$ are assigned to the left child of α, and processes with rank $\{i \geq \frac{q}{2}\}$ are assigned to the right child of α. At level $\log p$, each tree node is assigned to a single MPI process, which is also assigned the entire subtree of that node from level $\log p + 1$ to the leaf level (not displayed in Figure 1).

Some processes may own the skeletons and P matrix of the distributed node. We give an example in Figure 1; only $\{0\}$ in each node α stores the skeletons $\tilde{\alpha}$ and $P_{\tilde{\alpha}[\mathtt{lr}]}$. In the factorization we need to build matrices $U_\alpha, W_\alpha, Z_\alpha$, which in turn involve $P_{\tilde{\mathtt{l}}\tilde{\mathtt{l}}}$, $P_{\tilde{\mathtt{r}}\tilde{\mathtt{r}}}$, $K_{\mathtt{l}\tilde{\mathtt{r}}}$ and $K_{\mathtt{r}\tilde{\mathtt{l}}}$. These matrices are unnecessary in ASKIT but in INV-ASKIT we need to compute and store them in parallel. These computations involve collective broadcasts and reductions in the communicators of $\mathtt{l}, \mathtt{r}$ and point-to-point exchanges between the $\{0\}$ and $\{q/2\}$ ranks of the processes that were assigned to α. Finally matrix Z is stored at $\{0\}$. Figure 2 shows how these matrices are distributed for the matrix generated by the tree in Figure 1.

Algorithm II.3 `DistributedFactorize(`α`,`q`)`

1: **if** α is at level $\log p$ **then**
2: `Factorize(`α`).`
3: **else**
4: $\{< \frac{q}{2}\}$ `DistributedFactorize(`$\mathtt{l}, \frac{q}{2}$`).`
5: $\{\geq \frac{q}{2}\}$ `DistributedFactorize(`$\mathtt{r}, \frac{q}{2}$`).`
6: $\{0\}$ `Bcast` $P_{\tilde{\alpha}[l\tilde{r}]}$ inside the distributed node.
7: $\{0\}\rightarrow\{\frac{q}{2}\}$ `SendRecv` $\tilde{\mathtt{l}}$.
8: $\{\frac{q}{2}\}$ `Bcast` $\tilde{\mathtt{l}}$ using the $\mathtt{r}$ communicator.
9: $\{\frac{q}{2}\}\rightarrow\{0\}$ `SendRecv` $\tilde{\mathtt{r}}$.
10: $\{0\}$ `Bcast` $\tilde{\mathtt{r}}$ using the $\mathtt{l}$ communicator.
11: $\{< \frac{q}{2}\}$ Compute $K_{\mathtt{l}\tilde{\mathtt{r}}}$, $P_{\tilde{\mathtt{l}}\tilde{\mathtt{l}}}$ portions.
12: $\{\geq \frac{q}{2}\}$ Compute $K_{\mathtt{r}\tilde{\mathtt{l}}}$, $P_{\tilde{\mathtt{r}}\tilde{\mathtt{r}}}$ portions.
13: $\{< \frac{q}{2}\} W_{\mathtt{l}\tilde{\mathtt{r}}} =$ `DistributedSolve(`$\mathtt{l}, K_{\mathtt{l}\tilde{\mathtt{r}}}, \frac{q}{2}$`).`
14: $\{\geq \frac{q}{2}\} W_{\mathtt{r}\tilde{\mathtt{l}}} =$ `DistributedSolve(`$\mathtt{r}, K_{\mathtt{r}\tilde{\mathtt{l}}}, \frac{q}{2}$`).`
15: $\{0\}$ `Reduce` $P_{\tilde{\mathtt{l}}\tilde{\mathtt{l}}} W_{\mathtt{l}\tilde{\mathtt{r}}}$.
16: $\{\frac{q}{2}\}$ `Reduce` $P_{\tilde{\mathtt{r}}\tilde{\mathtt{r}}} W_{\mathtt{r}\tilde{\mathtt{l}}}$.
17: $\{\frac{q}{2}\}\rightarrow\{0\}$ `SendRecv` $P_{\tilde{\mathtt{r}}\tilde{\mathtt{r}}} W_{\mathtt{r}\tilde{\mathtt{l}}}$.
18: $\{0\}$ computes and LU factorizes Z.
19: **end if**

Factorization. In Algorithm II.3, we describe how (11) and (12) are computed. Each process uses its local MPI rank to make different decisions. E.g., if a statement is labeled with $\{< \frac{q}{2}\}$, then only processes with rank $< \frac{q}{2}$ execute it. Otherwise, all statements are executed by all processes. The whole communication pattern is simple. If a process owns the skeletons and P of α, then it sends P to all processes in the same node and sends its skeletons to all processes in its sibling tree node. In the pseudo-code, this is a little bit complicate, since we can't simply broadcast between communicators. Thus, some point-to-point communication is involved.

We will explain the algorithm using the yellow process (**Yellow**) in Figure 1 and Figure 2. At level 2, Algorithm II.1 is invoked, and the local tree of $\mathtt{r}(\alpha)$ is factorized. At level 1, **Yellow** returns from line 5 as rank $\{1\}$. At line 6, it receives $P_{\alpha[\widetilde{\mathtt{lr}}]}$, which will later be used to compute $P_{\alpha\widetilde{\alpha}}$ in the parent level. At line 7, it receives $\widetilde{\mathtt{l}}$ from **Green**, and at line 9 it sends $\widetilde{\mathtt{r}}$ it owns to **Green**. $\widetilde{\mathtt{l}}$ is used to evaluate $K_{\mathtt{r}\widetilde{\mathtt{l}}}$ at line 12. Evaluating $P_{\widetilde{\mathtt{r}}\mathtt{r}}$ doesn't require communication at this level, since all factors in the subtree are owned by **Yellow**. At line 14, $K_{\mathtt{r}\widetilde{\mathtt{l}}}$ is overwritten by $W_{\mathtt{r}\widetilde{\mathtt{l}}}$. Finally, **Yellow** sends its $P_{\widetilde{\mathtt{r}}\mathtt{r}}W_{\mathtt{r}\widetilde{\mathtt{l}}}$ to **Green**, and Z is computed and factorized on **Green**. At level-0, we only illustrate how $P_{\widetilde{\alpha}}[\widetilde{\mathtt{lr}}]$ that **Yellow** received at level 1 is used to compute the $P^{yellow}_{\widetilde{\alpha}\alpha}$ at line 11. By the definition of (12), the right portion of $P_{\widetilde{\alpha}\alpha} = [P^{green}_{\widetilde{\alpha}\alpha}, P^{yellow}_{\widetilde{\alpha}\alpha}]$ only requires $P_{\widetilde{\alpha}[\widetilde{\mathtt{lr}}]}$ and $P_{\widetilde{\mathtt{r}}\mathtt{r}}$. The first part $P_{\widetilde{\alpha}[\widetilde{\mathtt{lr}}]}$ is received at level 1, and the second part $P_{\widetilde{\mathtt{r}}\mathtt{r}}$ is computed at level 1.

Algorithm II.4 $w = \mathtt{DistributedSolve}(\alpha, u, q)$

1: **if** α is at level $\log p$ **then**
2: $u = \mathtt{Solve}(\alpha, u)$.
3: **else**
4: $\{< \frac{q}{2}\}$ $u = \mathtt{DistributedSolve}(\mathtt{l}, u, \frac{q}{2})$.
5: $\{\geq \frac{q}{2}\}$ $u = \mathtt{DistributedSolve}(\mathtt{r}, u, \frac{q}{2})$.
6: Compute $P_{\widetilde{\alpha}\alpha}u$ portion.
7: $\{0\}$ `Reduce` $P_{\widetilde{\alpha}\alpha}u$ portion and compute $ZP_{\widetilde{\alpha}\alpha}u$.
8: $\{0\}$ `Bcast` $ZP_{\widetilde{\alpha}\alpha}u$.
9: Compute $w = u - W_{\alpha\widetilde{\alpha}}ZP_{\widetilde{\alpha}\alpha}u$.
10: **end if**

Solver. Algorithm II.4 uses a similar bottom-up tree traversal. Algorithm II.2 is invoked when the traversal reaches level-$\log p$. At line 7, $\{0\}$ reduces all portions of $P_{\widetilde{\alpha}\alpha}u$ and computes $ZP_{\widetilde{\alpha}\alpha}u$ where Z was previously factorized. At line 8, $ZP_{\widetilde{\alpha}\alpha}u$ is broadcast by $\{0\}$. Finally, each process computes its own $w = u - W_{\alpha\widetilde{\alpha}}ZP_{\widetilde{\alpha}\alpha}u$.

Complexity. We use a network model to derive the parallel complexity. We assume that the network topology provides $\mathcal{O}(s \log p)$ complexity for `Bcast` and `Reduce` for a message of size s. For notational clarity we omit the latency cost.

Let $T^s(q)$ be the parallel complexity of Algorithm II.4. $T^s(1) = \mathcal{O}(n \log n)$ is the base case complexity of Algorithm II.2. We can derive the complexity recursively as $T^s(q) = T^s\left(\frac{q}{2}\right) + \mathcal{O}(ns + s^2) + \mathcal{O}(s \log q)$. Here $\mathcal{O}(ns + s^2)$ is the message transfer computation cost of (11), and $\mathcal{O}(s \log q)$ is the communication cost. The overall ***solve complexity*** with grain size $n = N/p$ is given by

$$T^s(n, p) = \mathcal{O}(n \log N)_{comp} + \mathcal{O}(s \log^2 p)_{comm}, \quad (15)$$

including both computation and communication. The number of messages required is $\mathcal{O}(\log^2 p)$ if one wants to include latency costs. We observe that the communication cost doesn't scale with the problem size N, but only with s and $\log^2 p$.

The complexity $T^f(q)$ of Algorithm II.3 is similarly given by $T^f(q) = T^f\left(\frac{q}{2}\right) + sT^s\left(\frac{q}{2}\right) + \mathcal{O}(ns) + \mathcal{O}(sd + s^2)$. Here $sT^s(\frac{q}{2})$ is the time in Algorithm II.4, $\mathcal{O}(ns)$ the time in (12), and $\mathcal{O}(sd + s^2)$ is the communication cost spent on broadcasting skeletons and P. Overall, the ***factorization complexity*** with grain size $n = N/p$ is

$$T^f(n, p) = \mathcal{O}(n \log^2 N)_{comp} + \mathcal{O}(s^2 \log^3 p + d \log p)_{comm}, \quad (16)$$

including both computation and communication. Observe that the sequential complexity is p times the parallel computation complexity, which indicates that the algorithm should be able to achieve almost linear scalability. The communication cost is minor since it does not depend on N.

Discussion. INV-ASKIT shares the common problems of all treecodes: (1) the parallelism diminishes exponentially when traversing to the root, and (2) the possible global synchronization during the level by level traversal. Solving (5) in a distributed node reduces the PW product to one process in the communicator. This takes the advantage of the collective `Bcast` and `Reduce`, but the parallelism diminishes. This efficiency degradation is not significant when p and s are small, but we will observe this effect in the scaling experiments. The load-balancing problem happens when the adaptive skeletonization is used to dynamically decide the rank of the off-diagonal blocks. The processes with the largest rank becomes the bottleneck. This load-balancing issues requires a dynamic scheduling and beyond the scope of this paper.

III. EXPERIMENTAL SETUP

We performed numerical experiments to examine the accuracy and the scalability of INV-ASKIT. Notice that the only alternative method that we could compare with INV-ASKIT for the datasets of the size and dimensionality is a global low-rank approximation of K. But in [24], we have already shown that for several bandwidths of practical interest the low-rank approximation is not a good scheme for approximating K.

We also applied INV-ASKIT to binary classification tasks on real-world datasets using kernel regression. Next we

Dataset	N_{train}	N_{test}	d	h
COVTYPE	500K	10K	54	0.07
SUSY	4.5M	10K	8	0.07
MNIST	1.6M	10K	784	0.10
NORMAL	1M-32M	-	64	0.31-0.17

Table II *Datasets used in the experiments. Here N_{train} denotes the size of the training set and corresponds to the N in* INV-ASKIT*; N_{test} is the number of test points used to check the quality of the classification. d is the dimensionality of points in the dataset and h is the bandwidth of the Gaussian kernel used in our experiments. We also performed tests with a polynomial kernel.*

summarize the hardware, datasets, and parameters used in our experiments. In the majority of our tests we use the Gaussian kernel function (1).

Implementation and hardware. All experiments were conducted on the Maverick and Stampede clusters at the Texas Advanced Computing Center. Maverick nodes have two 10-core Intel Xeon E5-2680 v2 (3.1GHz) processors and 256GB RAM. Stampede nodes have two 8-core E5-2680 (2.8GHz) processors and 32GB RAM. INV-ASKIT and ASKIT are written in `C++`, employing OpenMP for shared memory parallelism, Intel MKL for high performance linear algebra operations and Intel MPI for distributed memory parallelism. A custom direct "MatVec"[4] in ASKIT is optimized with `SSE2` and `AVX`. The classification code is written in `C++`, employing a Krylov subspace method (GMRES) from the PETSc library [4].

Iterative methods. Recall that INV-ASKIT can only compute $(\lambda I + \widetilde{K})^{-1}u$ without the sparse matrix S. While $S \neq 0$, ASKIT can better approximate K using nearest-neighbor pruning (ϵ_M will be smaller), but we can no longer compute its exact inverse directly. Still INV-ASKIT can be a good preconditioner. In this case, converging to the desired accuracy may take several GMRES iterations, but the classification accuracy may also be improved.

Datasets: In Table II, we give details of each dataset. In strong scaling and the classification tasks on Maverick, we use three real-world datasets: **SUSY** (high-energy physics) and **COVTYPE** (forest cartographic variables) from [21], **MNIST** (handwritten digit recognition) from [6]. In the weak scaling tasks, we generate a random 64D point cloud which are drawn from a 6D **NORMAL** distribution and embedded in 64D with additional noise. This resembles a dataset with a high ambient but relatively small intrinsic dimension, something typical in data analysis. We also report strong scaling results using a 16M **NORMAL** synthetic dataset.

Bandwidth selection and accuracy metrics. For the majority of experiments, we used the Gaussian kernel. The values are taken from [24], where they were selected using cross-validation with a kernel Bayes classifier. The regularization λ is also selected by cross-validation.

We test both versions of ASKIT (defined by the point-node separation criterion discussed in §II-B), which correspond to $S = 0$ and $S \neq 0$ in (9). We use three metrics to measure the quality of the approximations for the MATVEC and INVERT approximations of ASKIT and INV-ASKIT:[5]

$$\begin{aligned} \epsilon_M &= \|Kw - (\widetilde{K} + S)w\|_2 / \|Kw\|_2, \\ \epsilon_I &= \|w - (\lambda I + \widetilde{K})^{-1}(\lambda I + \widetilde{K})w\|_2 / \|w\|_2, \\ \rho &= \|u - (\lambda I + \widetilde{K} + S)w)\|_2 / \|u\|_2. \end{aligned} \tag{17}$$

The first error measures the ASKIT MATVEC error. If the off-diagonal blocks of K indeed admit low-rank, then ϵ_M should be small. Since evaluating Kw exactly ($\mathcal{O}(N^2)$ work) is impossible for large N, we estimate ϵ_M by sampling 1K points. We measure the condition of INV-ASKIT's INVERT for the case $S = 0$ using ϵ_I. If $(\lambda I + \widetilde{K})$ is well-conditioned, then ϵ_I should be small. We test $(\lambda I + \widetilde{K})^{-1}$ as a preconditioner for GMRES for solving $(\lambda I + \widetilde{K} + S)w = u$ and we report the normalized residual ρ upon termination of GMRES. Typically the desired ρ for the classification tasks is 1e-3. However, if the GMRES terminates while reaching the maximum iteration steps, then ρ will be larger. Overall, we measure the quality of a preconditioner according to the required iteration steps #iter and the achieved accuracy ρ.

Kernel regression for binary classification. Given a set of points with binary labels[6] , we split these points into the training and the testing sets. We use the training set to compute the weights for the regression. That is, given the correct classification labels u, where $u(i) \in \{-1, 1\}$ for each training point i, we solve $(\lambda I + \widetilde{K} + S)w = u$ for the regression weights w. Once we compute these weights, the label for a test point is given by $x \notin \mathcal{X}$ is $\text{sign}(\mathcal{K}(x, \mathcal{X})w)$. By comparing the true labels of the test points (which are given) to the computed ones using the kernel regression, we can estimate the classification accuracy as a percentage of correctly classified points. The classification errors are reported in the next section.

IV. EMPIRICAL RESULTS

We present three sets of results, for weak and strong scaling and for kernel regression. We conducted 29 experiments labeled from #1 to #29 in the tables. In Table III, we report numerical accuracy and strong scaling results. In Table IV, we report weak scaling for both Gaussian and polynomial kernels. Finally, we apply INV-ASKIT as a preconditioner in the binary classification problem. We show that INV-ASKIT's $(\lambda I + \widetilde{K})^{-1}$ speeds up the convergence rate up to $5\times$ and thus it is a good preconditioner for $(\lambda I + \widetilde{K} + S)$. We expect this to be the case when $\|S\|$ is sufficiently smaller

[4]See the General Stride Kernel Summation project: https://github.com/ChenhanYu/ks.

[5]Reminder on the notation: Here K is the exact kernel matrix and $\widetilde{K}$ is its ASKIT approximation with $S = 0$. By $(\lambda I + \widetilde{K})^{-1}$, we refer to the inverse based on the ASKIT approximation.

[6]The sets **MNIST** and **COVTYPE** have several classes. For those we perform a one-vs-all binary classification. Take **MNIST** dataset as an example, the two classes in the binary classification task are whether a point is the digit 3 or not.

#	core	T^M	T_c^f	T_f^f	T^s	T	η
COVTYPE			$h=0.07$, $\lambda=0.3$, $\kappa=2048$, $\tau=10^{-7}$, $\epsilon_M=1e-1$, $\epsilon_I=4e-12$.				
[redacted]	160	1.1	0.2	$80+64$	1.0	145	1.00
#2	320	0.6	0.4	$30+46$	0.8	77	0.94
#3	640	0.3	0.7	$10+31$	0.5	42	0.86
SUSY			$h=0.07$, $\lambda=10$, $\kappa=2048$, $\tau=10^{-3}$, $\epsilon_M=2e-1$, $\epsilon_I=3e-14$				
[redacted]	320	2.6	1.2	$636+666$	3.8	1307	1.00
#5	640	1.3	1.5	$269+409$	2.3	681	0.96
MNIST			$h=0.3$, $\lambda=0$, $\kappa=256$, $\tau=10^{-3}$, $\epsilon_M=6e-8$, $\epsilon_I=4e-13$				
[redacted]	160	2.5	0.1	$12.2+3.0$	0.26	16	1.00
#7	320	1.3	0.1	$6.6+2.1$	0.15	9	0.88
#8	640	0.5	0.1	$3.6+1.6$	0.08	5	0.80
NORMAL			$h=0.19$, $\lambda=0.1$, $\kappa=128$, $\epsilon_M=4e-1$, $\epsilon_I=4e-6$				
[redacted]	1,024	1.5	0.2	$37.6+57.3$	0.9	96	1.00
#10	2,048	0.8	0.3	$17.4+33.7$	0.5	52	0.92
#11	4,096	0.4	0.2	$8.6+21.0$	0.3	30	0.80

Table III *Strong scaling results. Experiments #1–#8 are done on Maverick using $m=2,048$ and $s_{max}=2,048$. Experiments #9–#11 are done on Stampede with $N=16M$, $m=512$ and $s=256$. T^M is the* MATVEC *time, T_f^f is computation time during the factorization, T_c^f is communication time during factorization, and T^s is solve time, all in seconds; η denotes the parallel efficiency.*

	Gaussian				Polynomial			
#	#12	#13	#14	#15	#16	#17	#18	#19
N	1M	4M	16M	32M	1M	4M	16M	32M
cores	20	80	320	640	20	80	320	640
h	0.31	0.24	0.19	0.17	-	-	-	-
ϵ_M	1e-1	2e-1	3e-1	3e-1	$\mathcal{O}(1e-12)$			
ϵ_I	1e-7	1e-7	6e-7	1e-6	6e-8	3e-7	1e-6	2e-6
T_c^f	0.0	0.1	0.1	0.2	0.0	0.0	0.1	0.1
T_f^f	295	396	509	574	300	401	509	574
T^s	2.5	3.1	3.8	4.2	2.5	3.1	3.8	4.2
T	298	399	513	578	303	404	513	578
η_f	0.3	1.1	4.2	8.2	0.3	1.1	4.3	8.2
η	1.00	0.92	0.88	0.85	1.00	0.92	0.89	0.85

Table IV *Weak Scaling experiment results on **NORMAL** synthetic datasets. All experiments use $m=512$, $s=512$, $\kappa=128$ and $\lambda=0.1$. The runtime is measured in seconds. η_f denotes* TFLOPS *and η denotes the relative efficiency.*

than $\widetilde{K}$. In the results the parameter τ indicates the accuracy tolerance used in the skeletonization phase of the ASKIT algorithm. s_{max} is the maximum skeleton size. κ is the number of the *nearest-neighbors*.

Accuracy. In Table III and Table IV, we report ASKIT's ϵ_M error for reference but our focus is on ϵ_I and ρ, which are related to INV-ASKIT. Note that ϵ_I is not always small due to the amplification of rounding errors by the conditioning of $(\lambda I+\widetilde{K})$. Both h and λ have strong impact on the condition number. For a Gaussian kernel matrix, a smaller h usually results in better conditioned systems since it approaches the identity matrix. Runs #6–#8 are examples where the systems are well-conditioned even with $\lambda=0$. With a larger h and lower d, $\tilde{K}$ becomes nearly rank-deficient. In #12–#19 we see that $\lambda=0.1$ is not sufficiently large, and that's why we observe large ϵ_I.

Scaling. Now, we discuss the weak and strong scaling results in Table III and Table IV. The overall INV-ASKIT time $T=T_c^f+T_f^f+T^s$ is divided into ***factorization time*** $T_f^c+T_f^f$ and the ***solving time*** T^s. The factorization time is split into (1) communication T_c^f and (2) computation T_f^f. T_f^f is further subdivided into local and distributed time. The communication part of T^s is negligible so we do not split T^s. Take #1 as an example, the communication time is 0.2 seconds, the local factorization time is 80 seconds, and the distributed factorization time is 64 seconds. The solving time is 1 second, and the overall time is 145 seconds.

In Table III we report the strong-scaling efficiency η. The base-line experiments are #1, #4, #6 and #9, which have $\eta=1.00$. Fixing the problem size, we measure the efficiency by doubling the core number each time. We observe that the communication cost T_c^f is insignificant; roughly speaking, doubling the number of cores will reduce the efficiency by 4%–12%. The efficiency degradation is due to diminishing parallelism during the computation of (5), where only the root rank in each subcommunicator works. We also measure the relative-to-peak FLOPS efficiency. The peak for 640-core on Maverick is 15.8[7] TFLOPS, and the peak of 4096-core on Stampede is 85.2 TFLOPS. The ratio FLOPS/peak from low-to-high is **MNIST2M** 8%, **NORMAL** 20%, **COVTYPE** 36% and **SUSY** 45%.

The scalability varies between datasets during the factorization, because there are load-balancing issues in the level-by-level traversal. Adaptive skeletonization results in different values of s for nodes in the same tree level. The node with the largest s becomes the critical path in each level. The load-balancing problem is especially serious in **MNIST** where the workload can vary by orders of magnitude. On the other hand, **SUSY** has the most balanced workload, which allows it to reach 45% peak floating point efficiency. We also report the MATVEC performance in the T^M column. The MATVEC in ASKIT is highly optimized and does not require any communication. Typically, it can reach near-linear scaling and more than 80% of peak floating point efficiency.

In Table IV, we report weak-scaling results. We perform experiments using Gaussian and polynomial kernels and we observe similar behavior. We report the FLOPS rate (η_f) and relative parallel efficiency (η). The base-line experiments are #12 and #16. The efficiency drops dramatically from 20-core runs to 80-core runs, since #12 and #16 do not involve MPI. From 80 cores to 640 cores, doubling the number of cores results only in 3% loss, which shows that the weak scalability is almost linear. The weak-scalability experiments can reach a higher efficiency, since there are no communication and load-balancing issues. As a result, INV-ASKIT can reach 52% theoretical peak floating point efficiency.

Classification. We conduct three classification experi-

[7]Each E5-2680 v2 processor can compute $3.1\times 8\times 10=248$ GFLOPS per second. The aggregate theoretical peak performance of 64 processors is $64\times 248=15,782$ GFLOPS ≈ 15.8 TFLOPS.

#	M	T_{train}	#iter	Train (%)	Test (%)	ρ
COVTYPE						$(\lambda I + \tilde{K})w = u.$
#20	I	231	100	99%	96%	4e-3
#21	$(\lambda I + \tilde{K})$	62	0	99%	96%	2e-11
COVTYPE						$(\lambda I + \tilde{K} + S)w = u.$
#22	I	1725	100	99%	97%	3e-3
#23	$(\lambda I + \tilde{K})$	950	43	99%	97%	1e-3
SUSY						$(\lambda I + \tilde{K})w = u.$
#24	I	3056	49	80%	79%	1e-3
#25	$(\lambda I + \tilde{K})$	870	0	80%	79%	1e-14
SUSY						$(\lambda I + \tilde{K} + S)w = u.$
#26	I	11519	47	80%	79%	1e-3
#27	$(\lambda I + \tilde{K})$	3419	7	80%	79%	1e-3
MNIST						$(\lambda I + \tilde{K})w = u.$
#28	I	185	5	100%	100%	1e-9
#29	$(\lambda I + \tilde{K})$	36	0	100%	100%	2e-14

Table V *Classification results using 640 cores. Other parameters are identical to the experiments in Table III. Here M denotes the preconditioner, which can be either I or $\tilde{K}$. The number if* GMRES *iterations is denoted by #iter. When it is* 0*, it means immediate convergence because* INV-ASKIT *is highly accurate. The GMRES solver will terminate while reaching 100 iteration steps or $\rho < 1e-3$.*

ments, following [24]. The training time T_{train} is presented in pairs to show both the runtime and the iteration steps of the GMRES solver. Four experiments are conducted for each dataset, combining different kernel matrix approximations ($\lambda I + \tilde{K}$ or $\lambda I + \tilde{K} + S$) and different preconditioners (I indicates no preconditioning or $\lambda I + \tilde{K}$). We present both training and testing (10K samples) accuracy, and we present the normalized residual ρ upon termination of GMRES.

Approximating the kernel matrix using $\tilde{K}$, we reach 96% accuracy in **COVTYPE**, 79% accuracy in **SUSY** and 100% accuracy in **MNIST**. Using the nearest-neighbor-based separation criterion which creates a nonzero S, we can reach a higher accuracy (97%) in **COVTYPE**. In #21, #25 and #29, we use $(\lambda I + \tilde{K})^{-1}u$ as initial guess, which allows GMRES to converge immediately to a high precision by only applying the preconditioner. The training time only involves the factorization time and the solving time, which are both deterministic. Without preconditioners, GMRES takes 3× to 5× longer to converge to the desired accuracy. During the cross-validation, time spent on training may be much longer depending on the combination h and λ.

In runs #23 and #27, we show that INV-ASKIT also works well as a preconditioner when we include the nearest-neighbors S in the ASKIT MATVEC approximation. Comparing #22 to #23, we can find that #23 with INV-ASKIT only takes 43 iterations to converge to 1e-3. On the other hand, #22 reaches the maximum iteration limit, only converging to 3e-3. Involving nearest-neighbors S can better approximate K in some cases. For example, #23 can reach a higher classification accuracy than #21. Overall INV-ASKIT provides a 2× to 3× speedup. For nonlinear classification methods like kernel logistic regression and support vector machines, the factorization can be amortized across several nonlinear iterations. In those cases the speedups will be even more dramatic.

V. CONCLUSIONS

We have introduced a new parallel fast factorization algorithm for kernel matrices defined on datasets in high dimensions. We evaluated our methods on both real-world and synthetic datasets with different kernel functions and parameters. We conducted analysis and experiments to study the complexity and the scalability of INV-ASKIT. These experiments include scaling up to 4,096 cores, calculations on very large data sets and achieving up 50% of peak FLOPS performance. We show that applying INV-ASKIT to classification problems can result in fast convergence and competitive accuracy. We do not assume symmetry, thus we automatically support variable bandwidth kernels. To our knowledge, INV-ASKIT is the only algorithm that can be used with high-dimensional data for a large variety of kernels and the only algorithm that supports shared and distributed memory parallelism.

Our future work will focus on further optimization of our implementation. In particular, there are opportunities to resolve the parallelism shrinking problem by dynamic scheduling. Also, we plan to explore other possible variants with lower factorization cost, extensions to GPU architectures, and integration with other higher-level algorithms like support vector machines.

ACKNOWLEDGMENT

This material is based upon work supported by AFOSR grants FA9550-12-10484 and FA9550-11-10339; by NSF grant CCF-1337393; by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Numbers DE-SC0010518 and DE-SC0009286; by NIH grant 10042242; by DARPA grant W911NF-115-2-0121; and by the Technische Universität München—Institute for Advanced Study, funded by the German Excellence Initiative (and the European Union Seventh Framework Programme under grant agreement 291763). Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the AFOSR, the DOE, the NIH, or the NSF. Computing time on the Texas Advanced Computing Centers Stampede system was provided by an allocation from TACC and the NSF.

References

[1] S. Ambikasaran, *Fast Algorithms for Dense Numerical Linear Algebra and Applications*, PhD thesis, Stanford University, 2013.

[2] S. Ambikasaran and E. Darve, *An $\mathcal{O}(n \log n)$ fast direct solver for partial hierarchically semi-separable matrices*, Journal of Scientific Computing, 57 (2013), pp. 477–501.

[3] S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O'Neil, *Fast direct methods for Gaussian processes and the analysis of NASA Kepler mission data*, arXiv preprint arXiv:1403.6015, (2014).

[4] S. Balay, M. F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, and H. Zhang, *PETSc Web page*. http://www.mcs.anl.gov/petsc, 2014.

[5] M. Bebendorf, *Hierarchical matrices*, Springer, 2008.

[6] C.-C. Chang and C.-J. Lin, *Libsvm: A library for support vector machines*, ACM Transactions on Intelligent Systems and Technology (TIST), 2 (2011), p. 27.

[7] J. Chen, L. Wang, and M. Anitescu, *A fast summation tree code for Matérn kernel*, SIAM Journal on Scientific Computing, 36 (2014), pp. A289–A309.

[8] A. Gittens and M. Mahoney, *Revisiting the Nystrom method for improved large-scale machine learning*, in Proceedings of the 30th International Conference on Machine Learning (ICML-13), 2013, pp. 567–575.

[9] A. Gray and A. Moore, *N-body problems in statistical learning*, Advances in neural information processing systems, (2001), pp. 521–527.

[10] L. Greengard, D. Gueyffier, P.-G. Martinsson, and V. Rokhlin, *Fast direct solvers for integral equations in complex three-dimensional domains*, Acta Numerica, 18 (2009), pp. 243–275.

[11] Greengard, L., *Fast Algorithms For Classical Physics*, Science, 265 (1994), pp. 909–914.

[12] W. Hackbusch, *A sparse matrix arithmetic based on H-matrices. part I: Introduction to H-matrices*, Computing, 62 (1999), pp. 89–108.

[13] W. Hackbusch, B. N. Khoromskij, and E. E. Tyrtyshnikov, *Approximate iterations for structured matrices*, Numerische Mathematik, 109 (2008), pp. 365–383.

[14] N. Halko, P.-G. Martinsson, and J. Tropp, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.

[15] K. L. Ho and L. Greengard, *A fast direct solver for structured linear systems by recursive skeletonization*, SIAM Journal on Scientific Computing, 34 (2012), pp. A2507–A2532.

[16] T. Hofmann, B. Schölkopf, and A. J. Smola, *Kernel methods in machine learning*, The annals of statistics, (2008), pp. 1171–1220.

[17] M. Izadi Khaleghabadi et al., *Parallel H-matrix arithmetic on distributed-memory systems*, (2012).

[18] R. Kriemann, *Parallel H-matrix arithmetics on shared memory systems*, Computing, 74 (2005), pp. 273–297.

[19] ———, *H-LU factorization on many-core systems*, Preprint, Max Planck Institute for Mathematics in the Sciences, 2014.

[20] S. Le Borne, L. Grasedyck, and R. Kriemann, *Domain-decomposition based H-LU preconditioners*, in Domain decomposition methods in science and engineering XVI, Springer, 2007, pp. 667–674.

[21] M. Lichman, *UCI machine learning repository*, 2013.

[22] W. B. March, B. Xiao, and G. Biros, *ASKIT: Approximate skeletonization kernel-independent treecode in high dimensions*, SIAM Journal on Scientific Computing, 37 (2015), pp. 1089–1110.

[23] W. B. March, B. Xiao, S. Tharakan, C. D. Yu, and G. Biros, *A kernel-independent FMM in general dimensions*, in Proceedings of SC15, The SCxy Conference series, Austin, Texas, November 2015, ACM/IEEE.

[24] ———, *Robust treecode approximation for kernel machines*, in Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Sydney, Australia, August 2015, pp. 1–10.

[25] W. B. March, B. Xiao, C. Yu, and G. Biros, *An algebraic parallel treecode in arbitrary dimensions*, in Proceedings of IPDPS 2015, 29th IEEE International Parallel and Distributed Computing Symposium, Hyderabad, India, May 2015.

[26] S. M. Omohundro, *Five balltree construction algorithms*, International Computer Science Institute Berkeley, 1989.

[27] C. E. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*, MIT Press, 2006.

[28] R. Schaback and H. Wendland, *Kernel techniques: From machine learning to meshless methods*, Acta Numerica, 15 (2006), pp. 543–639.

[29] S. Wang, X. S. Li, J. Xia, Y. Situ, and M. V. De Hoop, *Efficient scalable algorithms for solving dense linear systems with hierarchically semiseparable structures*, SIAM Journal on Scientific Computing, 35 (2013), pp. C519–C544.

[30] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*, Springer, 2004.

[31] C. Williams and M. Seeger, *Using the Nyström method to speed up kernel machines*, in Proceedings of the 14th Annual Conference on Neural Information Processing Systems, no. EPFL-CONF-161322, 2001, pp. 682–688.